

2pf+5

METHOD FOR MANAGING VIRTUAL MACHINES

The present invention refers to the domain of machines connected to a telecommunication networks, in particular the management of virtual machines.

The term virtual is very popular and generally refers to the modelling and simulation of a real system.

In the domain of computer machines, a classic example is the virtualisation of a storage system.

In replacement of direct access on a hard disk, a machine accesses a virtual disk of which, the machine does not know the physical specifications. The storage requests are processed by a suitable interface that receives the storage or recovery demands and carries out the conversion to the physical disk(s). Furthermore, the hardware stock can be heterogeneous, for example, units of different types (magnetic disk, Flash memory etc) but the virtual representation is unique for the final user.

The principle is simple, when the storage capacity on this virtual unit exceeds the physical capacity, new physical units are added in a transparent way for the user.

Companies such as VMWare®, IBM® or Connectix® propose solutions to provide software solutions towards a further step of the virtualisation, namely the creation of a complete virtual machine comprising the operational system part.

This solution allows to share a physical machine by creating several virtual machines, each complying with suitable criteria in terms of storage means, calculation means and communication means.

The applications running on this type of virtual machine have the advantage of being easily transportable because they are independent of a physical machine. The latter can develop with the advancement of technology without requiring modifications to the applications.

As explained above, if more important requirements are highlighted for a given application, the physical machine is replaced and then the application is reactivated without modification. The new display interface, network or others are processed in a specific way on the local layer of the machine and a unique virtual image is given with respect to the application by the virtual machine.

This configuration simplifies the maintenance and the testing of new applications. In fact, on the same machine, this application can operate transparently on several software environments.

This configuration has limits since it is only applied on one machine. In the network configuration, a server is initialized by this virtual machine and several computers can operate different applications that can each be attached to a suitable operational system (Linux, Windows® 95/98, Windows® NT...).

This virtuality presents limits because it does not take into account that one machine and only one application level is possible.

The aim of this invention is to propose a method allowing the use of several virtual machines including the applications and allowing the federation of these applications for the formation of services.

This aim is achieved through a method for managing applications, making use of at least two physical machines linked by communication means, these machines comprising a physical structure on which is loaded a first software layer adapted to this machine as well as a second layer forming a virtual layer on which virtual machines will operate, the latter comprising at least one application, this method comprising the following steps:

- connection of at least two physical machines for the formation of a physical machine network,
- loading of a control programme associated to the virtual layer of each physical machine,
- establishment of a dialogue between the control programme and a system management process,
- establishment of a physical machine network including the virtual layers,

- definition by the system management process of a service containing several applications,
- communication between the management process and each virtual layer in order to determine the status of the virtual machines associated to said virtual layer,
- assignation of a virtual machine to one of the virtual layers taking into account the characteristics of the application.

This method is thus based on the definition of a service that calls one or more applications. It is to be noted that a virtual machine can comprise one or several applications and that these virtual machines are managed by a physical machine disposing of a virtual layer compatible with the virtual machine of the application. The role of the management process is the control of each application in order to visualize the operation of a service containing several applications.

This management process can also determine if a service is active or available since it knows all the applications that forms said service. A service can, for example, be an e-commerce Internet site using a database application for the list of sale goods, an application for the presentation of information for the user interface part and an application for secure payment. Each of these applications, operating on a virtual machine, can be assigned to different virtual layers according to the requirements of these applications. The database application will require a large storage capacity while the presentation application will use rapid communication means.

The invention is placed at the level of application management, of optimization of resources for these applications and in a more general way, allows the development of these applications in terms of storage capacity, communication means and computer performance.

In fact, the management process carries out surveillance tasks of the applications and anticipates future requirements. In the case of the addition or replacement of a physical machine, the applications using this machine are redirected to other machines thanks to the management process. The latter

interrupts the application to displace the data stored pertaining to this application of the old machine towards the new machine. Once this copy process has been made, the application is assigned to this new machine and can be reactivated.

The invention will be better understood thanks to the following detailed description that refers to the enclosed drawings that are given as a non-limitative example, namely:

- Figure 1 describes the prior art for the operation of a virtual machine,
- Figure 2 describes the structure of the invention.

In Figure 1, the configuration such as that available at present is represented. A physical machine MPH disposes of an operating system OSB such as Linux, Windows® NT/2000. This operational system takes into account the physical configuration of the physical machine and manages the resources such as the storage unit HD. This system is specific to the material used, that is to say that it disposes of programmes (driver) pertaining to the material concerned (display, network card, hard disk etc.).

Once the operational system is installed, the virtual layer CV is loaded that will allow access to a virtual machine. From this step, it disposes of the capacity to operate generic machines for which, regardless of the physical configuration, the behaviour of these machines will be similar, in particular the interfaces towards the physical resources of the machine.

On this virtual layer, a virtual machine can be placed with each comprising one application APP1, APP2, APP3 using different operating systems (OS1 and OS2) according to the requirements. Access towards the virtual layer CV is carried out by means of a virtual material layer HDV. Taking the example of an access towards a serial port, the application APPL1 accesses a virtual serial port pilot that releases all the functionalities that can be expected of such a programme. This pilot transmits the instructions to the virtual layer that converts the instructions according to the serial port effectively installed.

This virtual machine configuration allows several applications to operate using different software environments.

Figure 2 shows the configuration of the invention in which two important innovations have been provided.

Associated to each virtual layer CV is a specific programme for process control MR that is responsible for managing the operation of the virtual machine(s) MV using this physical machine. This process control MR is in relation to a management process PG that centralizes all the data related to different machines.

This management process PG has the following functions:

- registration of the physical machines MPH, their configurations and their status,
- assignation of the applications APP on any of the virtual layers CV,
- supervision of operation of each virtual machine MV,
- stoppage of a virtual machine MV and displacement towards another virtual layer, this displacement being accompanied by the files pertaining to this virtual machine (database, programmes)
- service creation S1, S2, S3 depending on more than one application,
- supervision of service operation by the supervision of the applications APP or virtual machines MV forming said service.

One of the first tasks of the process PG is the inventory of the modules on a given network. This inventory registers on one hand the physical machinery MPH, their configuration in terms of storage, processing and communication resources, and on the other hand the virtual machines MV with their applications APP intended to function on one of the physical machines by means of the virtual layer CV.

Once the inventory has been carried out, the following step consists in defining the services containing several applications. The management process then assigns the different virtual machines to the physical machines, according to the requirements of each of these. The example of an e-commerce site is well suited to this kind of structure. This site is constituted by a first application APP1 of graphic interfacing with the user and the emphasis will be placed on the broadband connection capacity. A second application

APP2 manages the database of products and clients. It is characterized by large storage means and the third application APP3 is a secure payment application. These three applications form the e-commerce service S1 and must be operational in order for the site to be considered in order.

The next step is the reactivation of the virtual machines or more generally a service that will bring about the reactivation of several virtual machines.

The concept of a service containing several applications allows more detailed management. It is possible to associate the regulations on one hand on the applications comprising this service and on the other hand on the physical machines on which these applications function. It is to be noted that a service can request several identical applications that will have as operating conditions on at least two different physical machines.

One of the important conditions related to a service is the cue to start the different applications. In fact, certain applications require the presence of other applications to operate and consequently cannot reactivate with the latter.

In the same way that an application can be stopped, copied and reactivated, a service has the same properties. The copy will take into account the rules of the service and in the case that two physical machines are one of the working conditions, the copy process will take into account this rule and will ensure that the application A operates on a different machine to application B.

The management process PG also includes dynamic surveillance of the service. According to the example of an Internet site that disposes of at least two user interface applications A and B and of a database application, if the process PG certifies that the activity of the applications A and B exceeds a limit value, this process will reactivate a new application C of the same type.

In the case of the replacement of the machine MPH3 on which the application APP3 operates thanks to the virtual machine MV3, the process PG stops the virtual machine MV3 to assign it to the machine MPH2, for example. Before reactivating this virtual machine MV3, it is necessary for the payment data

stored in the storage unit of the machine MPH3 to be transferred to the machine MPH2.

It should be noted that this data is directly connected to the application APP3 and forms an inseparable unit. The organization and the division of this data block is the task of the application itself.

Once the copy of the data has been carried out, the virtual machine MV3 and its application APP3 can reactivate and the service S1 becomes available again.

According to the type of embodiment, the virtual layer CV and the operational base system OPB can only form one programme. The functionalities are brought together simply in one single product instead of two but both functions exist, namely a layer connected to the hardware towards the physical machine and a virtuality layer CV to offer to the external world a unique image.

Likewise, the specific programme for process control MR can be directly included in the virtual layer CV.

The management process PG also includes optimization tools. In fact, it is possible that the operating cost of one machine differs from another machine according to the technology used and the backup means (uninterruptible power supply, RAID system) and it is possible to determine the operating cost of an application and in the same way of a service.

Besides the physical characteristics of a machine, it is possible to take into account cost and security factors. This knowledge opens the way for the relocation of physical machines as represented by the machine MPH10. The dotted lines indicate that this machine is located in a different position to the other machines. Therefore, thanks to the configuration of the invention, it is possible to assign virtual machines not only to the machines connected locally but also to the machines physically installed on other sites according to the dotted line LB.

This possibility opens new prospects in terms of the management of computing equipment. In fact, external resources are in competition with internal resources. The only constraints are to be found at the level of communication means and the presence in the relocated machines of the same virtuality layer offered by local machines.

At the time of the configuration of an application, the attached parameters of the physical machine such as the cost per storage unit, the transferred packet or CPU time, as well as the security level are taken into account. It is possible at the time of the configuration of a virtual machine to prohibit the latter from being relocated, for security reasons for example.

The management process PG disposes of relation means in order to establish statistics regarding the use of the virtual machines and the applications. This tool allows the anticipation of the request and allows the development of the computing equipment before being faced with the overloading of the system.

In another embodiment, the virtual layer does not have as an aim the creation of a virtual machine but rather aims to allow the execution of a specific programme.

This is the case, for example, for machines with a Java™ 2 Platform Enterprise Edition (J2EE™) layer on which applications using a module Enterprise Java Bean (EJB) function. According to the invention, the process control programme MR is installed on the physical machine and dialogues with this execution layer J2EE™. Functionality is identical for virtual machines, namely that the management process PG registers the physical machines as well as the applications and establishes the catalogue of physical machines having an interpreter of a particular type and of the applications requiring this type of interpreter.

The role of the management process is thus to operate in a heterogeneous format in which physical machines are activated with virtual layers for the operation of virtual machines and other machines are activated according to particular environments such as J2EE™.

Once the registration has been carried out, the management process PG allocates the operation of an application according to the environment desired by this application.

According to one particular embodiment, illustrated by the arrow LA in Figure 2, the management process PG also dialogues directly with the virtual machine through a suitable control programme. Therefore, even before accessing the resources of a virtual layer CV, the process knows the status of the application that operates on this machine and can in this way access a higher quality control application. In fact, by directly accessing the virtual machine, it is possible to manage several applications on the same virtual machine. Although the virtual machine is assigned to a virtual layer of a physical machine, there may be several applications operating on this virtual machine. The requirements of each application on the interior of the same virtual machine as well as the statistics will be established by an application rather than by a virtual machine.